

# memcached: Distributed Caching in Python and Zope

presented to the  
Fredericksburg Zope / Python Users' Group  
2006/06/12

Tres Seaver  
Agendaless Consulting  
tseaver@agendaless.com

# **Demo: sharing session among appservers**

# memcached project

- Background
  - Built to support LiveJournal
  - Used by:
    - Slashdot
    - Wikipedia / MediaWiki
    - SourceForge
- Architecture
  - Primary use case: cache “expensive” SQL queries
  - “its a cache, not a storage!”
- Scalability
  - Cache distributes across multiple hosts

# Failover modes

- Rebalancing
  - keys get redistributed from “downed” host to remaining pool
  - thrashing when “downed” host returns
- Degraded mode
  - keys from “downed” host “miss” indefinitely
  - non-thrashing

# memcached server

- RAM based
- slab vs. malloc()-based allocation strategies
- expiry based on heuristics

# tugela

- Berkeley DB backend
  - stats results different
- expiry managed by separate daemon

# Client API

- client is proxy for a pool of servers
- key can be simple hashable, or else a tuple, '(hash, real\_key)'
  - hash selects the server from the pool
  - server weights used to allocate buckets

# Python Wrapper

- Wrapper marshalls arbitrary objects as pickles!
- Mapping API Methods
  - get / set
  - remove / delete
  - add / replace
  - incr / decr

# Use case: Shared Sessions

- Violation of “cache, not storage” principle
- Probably reasonable for “volatile” session data

# Use case: Shared RAM Cache Manager

- Closest to original design space of memcached
- Cache intermediate, expensive results
  - SQL queries
  - Page fragments

# Use case: Shared ZEO Client Cache

- Appservers query RAM, then memcache, before loading object from ZEO storage
- Reduce load on storage at appserver start / restart

# Resources

- memcached project homepage:  
<http://www.danga.com/memcached/>
- python-memcache client bindings,  
maintained by Sean Reifschneider at  
Tummy.com:  
[http://www.python.org/pypi/memcached/1.2\\_tummy6](http://www.python.org/pypi/memcached/1.2_tummy6)
- tugela project homepage:  
[http://meta.wikimedia.org/wiki/Tugela\\_Cache](http://meta.wikimedia.org/wiki/Tugela_Cache)
- mcdutils product homepage:  
<http://agendaless.com/Members/tseaver/software/mcdutils>

# Questions

- Tres Seaver, [tseaver@agendaless.com](mailto:tseaver@agendaless.com)